

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

METHOD FOR WIRELESS DOWNLINK SCHEDULING

Background of Invention

[0001] The present invention relates to wireless data networks and, more particularly, to scheduling of downlink data traffic in a multirate multi-code wireless data network.

[0002] In wireless communications systems, numerous techniques are known to isolate communication links from one another. The conventional advanced mobile phone system (AMPS) in North America uses frequency division multiple access (FDMA) which assigns each call within a cell a unique pair of RF channels. With regard to IS-136 digital extensions to AMPS, time division multiple access (TDMA) is used to divide the RF channel into frames defining a predetermined number of time slots, thereby increasing the number of calls that can be handled on a single channel. In code division multiple access (CDMA) systems, all of the users share the entire transmission bandwidth and users are distinguished by the use of signatures – referred to in the art as "codes" – assigned to them. A CDMA receiver will decode only the call that is modulated with the code that uniquely identifies the proper call. A more recent air-interface being actively studied is based on wideband orthogonal frequency division multiplexing (OFDM), where the wideband channel is divided into narrow frequency tones, but the transmission frame is assembled using a Fast Fourier Transform (FFT).

[0003] Recent advances in next-generation "3G" wireless communication systems envision advanced architectures with highly improved data transmission rates. For example, the two main CDMA proposals competing for next generation voice/data standards, namely CDMA2000 and wideband CDMA (WCDMA), both envisage higher rates through assigning multiple codes to users (referred to in the art as "code aggregation") and variable rates through coding techniques. For example, WCDMA

allows for the following discrete set of data rates {9.6, 19.2, 38.4, 76.8, 153.6, 307.2, 614.4 kbps}. Similarly, the High Data Rate (HDR) system designed by Qualcomm for a purely data system, uses a set of 16 orthogonal codes occupying a physical bandwidth of 1.2288 MHz that only allows a certain discrete set of data rates (i.e. 38.4, 76.8, 102.6, 153.6, 204.8, 307.2, 614.4, 921.6, 1228.8, 1843.2, 2457.6 kbps over a time frame of length 1.67 milliseconds). These discrete rates translate to discrete allowable power assignments. There is a pilot signal in the broadcast control channel that enables the access terminal to measure the link channel conditions and this is reported back to the base station, permitting it to assess the transmission conditions for users at a time-scale of a few milliseconds per measurement. Similar physical, networking, and systems issues are addressed in CDMA2000 and WCDMA systems.

[0004] Providing fine-grained quality of service to users in such next-generation wireless networks gives rise to a new class of scheduling problems. Providing consistent quality of service to mobile users in the downlink requires management of both power (rate) and codes for every user. Wireline scheduling and resource allocation algorithms cannot be directly applied to manage the downlink. Wireless networks have unique characteristics, such as location dependent channel errors. In wireless systems, channels have variable attenuation depending on the geographical location of the user. This is mainly due to multipath impairments and radio propagation losses. Unlike traditional scheduling scenarios, in a wireless environment, the scheduler must consider channel state to provide reasonable quality of service. Random channel errors also result in performance problems for transport protocols such as TCP, since TCP typically interprets channel error as errors caused due to congestion. Link layer retransmissions do not help, but aggravate the situation since they interfere with TCP's *rtt* computations. Most solutions to this problem propose intercepting the connection at the base station and creating two logical connections, like a proxy. As a proxy, the base station has access to information such as the request sizes. In order to manage the link more effectively, the base station must make use of this information, all while using state information obtained from MAC layers judiciously in resource scheduling.

Summary of Invention

[0005] The present invention is directed to scheduling of data transmissions in a multi-code (or multi-channel) wireless data network that supports multiple data rates. In accordance with an aspect of the invention, the overall power(rate) and the number of codes to assign to a data request over an entire schedule are selected based on system parameters and other information such as the size of the data request and transmission characteristics to the particular receiver of the data request. The inventors have found it advantageous to utilize resource augmented competitive analysis in selecting the choice of overall power and number of codes. The inventors have also found it advantageous to select the overall power(rate) and the number of codes without regard to the discrete nature of the data rates in the wireless data network. Instead, the results are rounded so that every selected code is assigned a power that achieves a data rate feasible for the particular receiver and system. Then, the modified power (rate) and number of codes is allocated in each scheduling frame in accordance with a quality of service metric. For example, and without limitation, the system can attempt to minimize maximum response time by allocating the requests in order of the time of release. Alternatively, the system can attempt to minimize a total weighted response time.

[0006] In accordance with another aspect of the invention, where the handling of a request leaves power/codes unused in a particular time slot, then other jobs can be "packed" into the time slot. For example, requests with the earliest release time can be given higher priority over other job requests, but other job requests can be considered in non-decreasing order of their release time where there are unused power/codes in that time slot.

[0007] The present invention provides a practical online method for scheduling the data transmissions, despite the inherent complexities of the scheduling problem. These and other advantages of the invention will be apparent to those of ordinary skill in the art by reference to the following detailed description and the accompanying drawings.

Brief Description of Drawings

[0008] FIG. 1 illustrates a wireless data transmission system appropriate for use with the present invention.

- [0009] FIG. 2 is a graph illustrating an instance of data rates as a function of SINR.
- [0010] FIG. 3 is a flow chart of processing performed by a wireless network, in accordance with an aspect of the invention.
- [0011] FIG. 4 is a flow chart illustrating an embodiment of a 2D-FIFO scheduling.
- [0012] FIG. 5A and 5B is a flow chart illustrating an embodiment of scheduling using a general selection function.

Detailed Description

- [0013] FIG. 1 illustrates a wireless data transmission system appropriate for use with the present invention. Each cell 110 in a packet cellular architecture has a base station 120 which transmits data to one or more receivers 111, 112, ... 113, which may be mobile or fixed. The data transmission system is assumed to have the capability to transmit at multiple data rates and to assign multiple codes or channels to a single receiver. Herein and in the claims, the inventors use the term "codes" not with specific reference to codes in CDMA but generically with reference to either codes or channels, thereby encompassing the general capability of assigning either codes or channels to a single receiver. The receivers/hosts 111, ... 113 may constitute any of a number of packet data-enabled wireless communication equipment, including, without limitation, cell phones, personal digital assistants, computers equipped with wireless modems, etc. The base station 120 is typically connected to a mobile switching center ("MSC") 130 which manages the scheduling of data transmission to the receivers/hosts 111 ... 113. The MSC 130 provides access to packet-switched data network 100. Network 100 can be any type of data network, although in FIG. 1 network 100 is depicted as being based on the TCP/IP protocol suite, such as the Internet. Network nodes 101, 102, e.g. routers, forward packets to and from other hosts 140 connected to the packet-switched data network 100. FIG. 1 shows for illustration only the protocol layers for a conventional packet data application using the Hyper Text Transfer Protocol (HTTP).
- [0014] The MSC 130 and the base station 120 are responsible for handling and scheduling all data requests to and from the packet data hosts in the cell 110, i.e. handling both uplink (from mobile users) and downlink (to mobile users) requests.

Both uplink and downlink channel performance are crucial for the overall system. The focus of the analysis herein is on the downlink and non-real time traffic, e.g. data browsing, downloads, etc. Downlink traffic is expected to dominate over time and data traffic typically tends to have asymmetrically large downlink demand.

[0015] It is advantageous to abstract the scheduling problem as follows. The base station 120 has a total power P to transmit. For analysis purposes, time is assumed to be partitioned into equal width windows referred to as "time frames", the width being Δ . It is also assumed that there are a total of C codes that can be assigned to users in a time frame. The communication channel is modeled as follows. Where the base station 120 transmits power p_j to a user j , the signal-to-interference-plus-noise ratio (SINR) is given by $\text{SINR} = (g'_j p_j) / \sigma^2$, where g'_j is a scalar factor referred to herein as the "physical" gain which reflects the physical channel attenuations of the user j and where σ^2 is the total noise power including interference. As is known in the art, SINR determines the rate of transmission of packets to the user as well as the probability of errors in transmission. The data transmission rate $r(\cdot)$ as a function of the SINR can be represented as a concave logarithmic function:

$$r(x) = \Delta \text{rbps}(x) = \Delta W' \log \left(1 + \frac{x}{\Gamma} \right)$$

where $\text{rbps}(\cdot)$ is the rate in bits per second, $r(\cdot)$ is the bits received per time frame, Γ is dependent on the coding gain from the physical layer error-correcting code, and W' is the spectral bandwidth used. Both Γ and W' are system parameters, which are assumed to be constant. FIG. 2 shows an instance of data rates as a function of SINR from both reported measurements and from the above equation. Note that the rate function is not linear in its argument and is, in fact, concave. Note also that the rate function already embodies the effect of variable rate error-correcting coding schemes in the physical layer, as is typical in next generational wireless systems. For notational convenience, $W = \Delta W'$ and $g_j = g'_j / \Gamma \sigma^2$ denotes the "channel" gain as compared to the physical gain.

[0016]

Requests for data transmission arrive in the system over time at the beginning of time frames. The i th request is referred to by the time frame of its arrival time a_i (time and time frame are referred to herein interchangeably when no confusion arises). At the time of the arrival of a data request, the size of the data request s_i

(measured, for example, in bits) and the channel gain g_i of the user who made the request is assumed to be known. It is also assumed for simplicity that the channel conditions of the users are constant over the scheduling period. Given the rate function shown in FIG. 2, the following useful properties of the communication channel can be shown. In a continuous power (rate) case, it can be shown that if $c \geq 1$ codes are assigned to a user u , then it is optimal to divide the total power p allocated to that user equally among the codes assigned. Also, due to the concavity of the rate function, it can be shown for the discrete rate case that, when the power on a code is cut in half, the rate is cut by at most a factor of four (this is instead of a lower bound factor of two in the continuous case).

[0017] The scheduling problem is to determine an assignment of power and codes to each user in each time frame so that the following conditions hold:

- (a) All requests get the requested data size.
- (b) Only a discrete set of data rates (or equivalently, minimum power per discrete rate) is allowed. These rates are denoted $R(1)$, $R(2)$, ... and have the property that $R(i)/R(i-1) \leq 2$. This relationship holds for existing next generation wireless data system proposals like HDR. Nevertheless, the results are not limited to a factor of 2; where the discrete rates are more spread out, but bounded by some constant, all of the results will apply with minor changes in the claimed bounds.
- (c) Various quality of service metrics are optimized. Common metrics of performance comprise, without limitation, minimizing the maximum response time; minimizing the total weighted response time; etc.

It is assumed that requests may be served over several time frames with different sets of codes at each time frame, i.e. the requests may be preempted or migrated.

[0018] It can be shown that the scheduling problem is NP-complete, even in very restricted cases, which means that efficient polynomial time algorithms exist only if $P = NP$, a complexity-theoretic outcome that is unlikely to be true. Nevertheless, in accordance with an embodiment of the present invention, it is possible to design an online scheduling methodology that advantageously makes all of the scheduling decisions at any instant based on only knowing requests released prior to that instant and without any assumption on requests that will be presented in the future. The online

algorithms optimize the quality of service metrics of interest by relying on the unusual technique of resource augmented competitive analysis.

[0019] In resource augmentation analysis, an optimum found by an ideal adversary is compared with the value of the solution found by an algorithm that has more resources than the adversary. See, e.g., B. Kalyanasundaram and K. Pruhs, "Speed is as powerful as clairvoyance," in IEEE Symposium on Foundations of Computer Sci., pp. 214–221, 1995. The ideal adversary knows the entire input instance in advance and serves it optimally; there is no bound on the time the adversary needs to identify the optimal solution. Formally, the algorithm for the scheduling problem is said to be an $(\alpha, \beta, \gamma, \delta)$ approximation if it provides a β approximation of the optimum when the sizes of the user requests is scaled down by a factor α and the number of codes (the power) used by the algorithm is at most γ (δ , respectively) times the number of codes (power, respectively) used by the optimum solution for the original input. Resource augmented analysis provides a theoretical way of understanding the inherent structure of the scheduling problem. It also provides a framework for a systems manager to analyze and properly provision for a concrete quality of service.

[0020] Given an $(\alpha, \beta, \gamma, \delta)$ approximation algorithm A , a $(1, \beta, \alpha\gamma, \alpha\delta)$ approximation algorithm A' can be obtained as follows: first apply A to the given instance and let X be the obtained solution that allocates power and codes to users. For each frame $x \in X$, A' uses α copies of x and allocates these frames in the same way as x . Clearly, the set of new frames allows the system to answer α times the demand satisfied by x . Note that A' uses α times more codes and total power than A . This implies that all of the results stated for algorithms working on requests of reduced size can be transformed into results for algorithms working on the original instance of the problem at the expenses of some extra codes and power allocated to the system. Accordingly, α need not be directly forced to 1 in the analysis below since no generality is lost in the process.

[0021] FIG. 3 is a flowchart of processing performed by the wireless network, in accordance with an aspect of the invention. At step 301, an allocation of codes and power is determined for a "continuous" rate case that satisfies a resource-augmented demand over the entire schedule. It should be noted that the inventors have found

experimentally that over-provisioning of codes is not as effective as over-provisioning of power when enhancing performance using resource augmentation. The rate (power) is allowed to take on any value in a continuous range rather than be restricted to a discrete set. The total transmission power is evenly split among the codes assigned to the request for the reduced demand. This may correspond for that specific user to a non-feasible transmission rate at the receiver. To move from the continuous case to the discrete case, it is necessary to round the power assignment to a value that sustains one of the discrete transmission rates. Accordingly, at step 302, a "rounding" procedure is applied that assigns to every code a power that achieves a feasible rate at the receiver side. The rounding scheme results in a new estimation of the overall number of codes that are required to complete the transmission. Given this total number and the power assignments to each code, it is possible then to determine the minimum number of frames that are necessary to complete a request in case it receives the entire set of codes and power available in each frame. Finally, at step 303, a selection function is invoked to allocate the codes and powers to users, frame by frame. The selection function is invoked when a scheduling job is released or when a job is completed. By choosing different selection functions, as further described below, different scheduling strategies are implemented.

[0022]

FIFO(P/C) AND 2D PACKING: FIG. 4 is a more detailed flowchart of processing performed by a wireless network, in accordance with an advantageous embodiment of the inventors' online heuristics. It is well known in the processor scheduling literature that the online algorithm Earliest Release Time (ERT) or First In First Out (FIFO) is optimal for minimizing the maximum response time on a single machine and is a 3 approximation algorithm on parallel machines. FIG. 4 illustrates an on-line scheme where each user is given a power P/C per code and is served by a FIFO scheduling discipline. It can be shown that this can achieve optimum results if every request is reduced to 50% of its original size. Note that this translates to augmenting power and codes by at most a factor of 2 for satisfying the full demand. Accordingly, although the approach of the present invention is described below in terms of a reduced request size, this is analytically interchangeable with augmenting the resources required to provide for the full request. In other words, the reduction in demand translates to satisfying the entire demand but on a resource augmentation factor

equal to the reduction factor.

[0023] Accordingly, at step 401, a reduced demand s_j^r is assigned, e.g. a reduced demand that is 50% of the original demand s_j . Satisfaction of the demand s_j requires that $\sum_t R_j = s_j$, where R_j is the sum of all transmission data rates over the set of codes assigned to user j in time frame t . At step 402, a number of codes k_j^r is assigned with a power allocation of P/C until the reduced demand is met. The number of codes k_j^r represents the number of codes allocated to user j that allows satisfaction of reduced demand s_j^r according to the gain factor of user j . Steps 401–402 correspond to step 301 in FIG. 3. Steps 403–404 represent an implementation of the rounding scheme described above with reference to step 302 in FIG. 3. The assignment of a power P/C to every code may correspond to a non-feasible transmission rate at the receiver for some specific user. To move from the continuous to the discrete case, it is necessary to round the power assignment to a value that sustains one of the discrete transmission rates. As shown in FIG. 4, two different kinds of rounding of a power z assigned to a code are performed. At step 403, if there exists a power $z_1 \in (z, 2z]$ corresponding to a discrete rate, then that power z_1 is assigned to the code. At step 404, on the other hand, if there exists a power $z_2 \leq z$ corresponding to a discrete rate, then that power z_2 is assigned to the code. For each user, a rounding is chosen that gives the higher rate, if the user were given all of the resources, i.e. all the power and codes. Finally, as further described below, the jobs are scheduled in accordance with an advantageous selection function. For example, where it is desired to minimize the maximum response time, it is advantageous to schedule the jobs in order of release time, in accordance with the FIFO scheduling discipline. When a job is completed, the pending user request is selected, if any, with the earliest release time. It can be shown such an allocation scheme satisfies all of the users demands.

[0024] Since the rounding of the rates shown in FIG. 4 might result in some power and/or codes being unused in a slot, it is desirable to design discrete-rate online methodologies that minimize this potential waste of resources in order to reduce the maximum response time. Three advantageous examples of such methodologies are described below. Given a job the power per codes corresponding to the discrete bit rate is the same for all of these embodiments. They differ only in the way the jobs are

selected for receiving service.

1. FIFO: this is an application of the traditional FIFO algorithm. The request currently in the system that has the earliest release time is always selected.
2. 2D-FIFO: the request currently in the system that has the earliest release time has higher priority over other job requests. However, if this job leaves power/codes unused in that time-slot, other jobs in the system are considered in non-decreasing order of their release times. This embodiment is illustrated in FIG. 4 at steps 405-406.
3. 2D-PIKI: the request currently in the system that has the highest value of power per code is selected. If this job leaves power/codes unused in that time-slot, other jobs in the system are considered in non-increasing order of the power per code. This scheme aims to achieve a better packing in each time-slot, in order to reduce the completion time.

Due to the discrete nature of the rate set, in certain slots the scheduler may have some codes and some power that cannot be assigned to any job in the system. In such a situation, the scheduler will choose the first job that received service in the slot and give it the best possible discrete rate with the remaining power and codes. This is applicable to all of the three embodiments described above. Note that no methodology guarantees that all the power and codes will be used in every slot.

[0025]

GENERAL SELECTION RULE: FIG. 5A and 5B illustrate an alternative allocation and rounding scheme, in accordance with another embodiment of FIG. 3. A general selection rule is applied at release time to determine the overall power and the number of codes to assign to a request over the entire schedule. The general selection rule selects a pair p, k that minimizes:

$$\frac{p_j^c}{P} + \frac{k_j^c}{C}$$

where p_j^c is the total feasible transmission power and k_j^c is the number of codes allocated to user j that allows satisfaction of demand s_j according to the gain factor of user j . Satisfaction of the demand s_j requires that $\sum_t R_j = s_j$, where R_j is the sum of all transmission data rates over the set of codes assigned to user j in time frame t . The possible choices, although allowed to range in a continuous manner, are nevertheless restricted in a way to allow a transformation for the discrete case. Denote by x^{\min} and x^{\max} , respectively, the minimum and maximum power per code that

allows a feasible transmission rate to user j . The maximum power per code is at most the maximum power that can be allocated in a frame, namely $x_j^{\max} \leq P$. These constraints lead to consideration of only those pairs such that $x_j^{\min} \leq p_j^c / k_j^c \leq x_j^{\max}$. It should be noted that the selection rule may assign to a user more power and codes than is available in a single frame, for which a user will in general receive codes from several different frames.

[0026]

The values selected by an algorithm according to the rule above are denoted p_j^m and k_j^m . Analogously, the values selected by an ideal adversary that has complete knowledge and that finds an optimal solution for the given objective function are denoted p_j^{OPT} and k_j^{OPT} . The above algorithm solution is related to the optimal solution in the case in which the quality of service provided to every user is slightly degraded. The reduced demand per user is denoted s_j^r and the overall power and number of codes allocated by the general selection rule when applied to the reduced demand are denoted p_j^r and k_j^r , respectively. Assume that the demand is reduced to a fraction $1/(4(1+\epsilon)^2)$, $\epsilon \geq 0$, of the original demand s_j , as set forth in step 501 in FIG. 5 (recall that this reduction in demand is equivalent to satisfying the entire demand with a resource augmentation factor equal to the reduction factor). This implies:

$$\begin{aligned} \frac{p_j^r}{P} + \frac{k_j^r}{C} &\leq \frac{1}{2(1+\epsilon)} \left(\frac{p_j^m}{P} + \frac{k_j^m}{C} \right) \\ &\leq \frac{1}{2(1+\epsilon)} \left(\frac{p_j^{\text{OPT}}}{P} + \frac{k_j^{\text{OPT}}}{C} \right) \\ &\leq \frac{1}{1+\epsilon} \max \left\{ \frac{p_j^{\text{OPT}}}{P}, \frac{k_j^{\text{OPT}}}{C} \right\}. \end{aligned}$$

For the equation above, consider the general selection rule when applied to the original demands. It selects the pair p_j^m, k_j^m that minimizes the earlier expression above. Since the allocated power is equally split among the codes, by reducing the number of codes for a factor $1/(2(1+\epsilon))$, a fraction $1/(2(1+\epsilon))$ of the demand is satisfied. Moreover, by reducing by a factor $1/(2(1+\epsilon))$ the total allocated power, a fraction $1/(4(1+\epsilon)^2)$ of the original demand is still satisfied. Accordingly, at step 502, the above general selection rule is applied to the demand to obtain a power, code solution that is augmented (or, equivalently, a solution that is applied to a demand that is reduced). It should be noted that this solution has allocated power per codes that is at least $x_j^{\min} / (2(1+\epsilon))$ and at most x_j^{\max} . The ratio p_j^r / x_j^r will, therefore, be allowed to be in the range $[x_j^{\min} / (2(1+\epsilon)), x_j^{\max}]$ in order to

apply the relation given by the equation above.

[0027] It should be noted again that although FIG. 5 (and FIG. 4) refer to reducing the demand, as explained above, a reduction in demand translates to satisfying the entire demand using a resource augmentation factor equal to the reduction factor. The two methods are interchangeable. By increasing the amount of resources servicing the request by a proportionate factor, the entire demand may be satisfied.

[0028] Steps 503, 504, and 505 in FIG. 5A illustrate an implementation of the rounding scheme. The transformation maintains the approximation of the algorithms in the continuous case at the expense of some extra codes and some extra power allocated in each frame. Let x_j^{\min} and x_j^{\max} be the minimum and maximum power corresponding to a transmission rate for user j in the discrete case. It is assumed that for every $x \in [x_j^{\min}, x_j^{\max}]$ there exists a power assignment $\bar{x} \in [x/2, 2x]$ corresponding to a discrete transmission rate. Three different kinds of rounding of a power z assigned to a code are performed, as set forth in FIG. 5A. At step 503, the power is rounded up to a discrete rate where a corresponding power exists between z and $2z$. At step 504, the power is rounded down to a discrete rate where a corresponding power exists between $z/2$ and z . At step 505, the power is rounded up to x_j^{\min} if z is less than x_j^{\min} .

[0029]

The rounding scheme illustrated by FIG. 5 can be utilized for codes with allocated power bigger than $x_j^{\min} / (2 + \epsilon)$. For certain selection functions, however, as further described below, a code assigned to a user j can have an allocated power less than this amount if either (a) it is the last assigned code in a frame with power y , say c_1 , or (b) it is the first assigned code in the successive frame, say c_2 , allocated with power $x - y$. In this case, the following rounding scheme can be used: (1) If $x < x_j^{\min}$, then round up code c_1 to x_j^{\min} and assign power 0 to code c_2 . (2) If ($x > x_j^{\min}$ and $y < x_j^{\min} / (2 + \epsilon)$) then round up to x_j^{\min} the power allocated to c_1 ; else (3) If ($x > x_j^{\min}$ and $y \geq x_j^{\min} / (2 + \epsilon)$) then insert a code with power x_j^{\min} in the previous frame. It can be shown formally that the above allocation scheme for the discrete case satisfies all user demands. Moreover, it can be shown that algorithms for the continuous case working on demands reduced by a factor $1/(4(1 + \epsilon)^2)$ can be

turned into an algorithm for the discrete case by allocating power at most $(3 + 2\epsilon)P$ and a number of codes at most $2C + 1$ in every frame (again, this is equivalent to satisfying the entire demand by augmenting the resources by an identical factor).

[0030] Finally, steps 506, 507, and 508 in FIG. 5B illustrate another implementation of the FIFO/Earliest Release Time (ERT) selection function. At step 506, when a job is completed, the user request with the earliest release time that has not been completed yet is selected for scheduling. Once the job is selected to be scheduled, a total transmission power is evenly split among the codes assigned to the request. At step 507, the codes and power are allocated frame by frame. A user j is allocated starting from the first frame t where some codes and transmission power are still available. At step 508, the allocation of the user request proceeds until one of the indicated three events occurs, i.e. (1) the codes and transmission power have all been assigned to the user; (2) the remaining total transmission power has already been assigned in the current frame and a portion of the power needs to be allocated to the next frame; or (3) all the codes have been assigned in the frame and codes from the next frame must be used. As illustrated by the embodiments above, it can be shown that this selection policy proves to be a good heuristic for minimizing the maximum response time.

[0031] SELECTION FUNCTION: As alluded to above, the allocation of the power and codes should be performed in a manner consistent with the quality of service metric to be optimized. For example, and without limitation, common metrics of performance are to (a) minimize the maximum response time, where response time (sometimes referred to as "flow" in the art) for request i is $c_i - a_i$ if request i completed by time c_i ; or (b) minimize total weighted response time, $\sum_i w_i (c_i - a_i)$, where arbitrary weights w_i are specified for each request i . If the weight for each job were 1, then this becomes the traditional average response time measure. If the weight is proportional to $1/t_i$ where t_i is the time it takes to service the i th request in a completely unloaded system, that is, when all codes and power assigned to request i . This measure, $(c_i - a_i) / t_i$, is known in the art as the "stretch" of a job. The response time metric is skewed towards large jobs since jobs with large service times also tend to have large response time. On the other hand, the relative response metric is independent of size making it more fair to all job classes. Since data requests in the

emerging data systems and applications would very likely be heterogeneous, relative response could be an attractive metric. Other weight functions may also be useful, although the two above are most common and will therefore be the focus of discussion.

[0032] It is straightforward to show that different allocation schemes for the same input data will be optimal, depending on the chosen quality of service metrics. For example, given two requests with significantly different gains, one can readily show that a scheme that allocates all codes and power to a request is not optimal when the quality of service measure is to minimize the maximum stretch. But the same allocation scheme can be optimal given different input data. The following are some different examples of advantageous selection functions with useful properties:

[0033] A. *Earliest Release Time (ERT)*. This selection function schedules jobs in order of the time of release, i.e. the request k currently in the system having minimum $a(k)$ is selected. As discussed above, it can be shown that this selection policy proves to be a good heuristic for minimizing the maximum response time.

[0034] B. *Highest Density First (HDF)*. This selection function schedules at any time t a request not completed yet with maximum ratio w_j / s_j , where w_j is a weight associated with job j . The scheduling of a user request is preempted in a frame when a user with higher density is released. When a user request is completed, that request with highest density is scheduled. The HDF selection function is advantageous for optimizing the weighted response time problem. It can be shown that the HDF selection function, for any $\epsilon > 0$, is a $(1 + \epsilon)/\epsilon$ approximation for minimizing the weighted flow time if every request is guaranteed for a fraction $1/(4(1 + \epsilon)^2)$ of the original demand.

[0035] C. *Shortest Processing Time (SPT)*. This selection function chooses the job k having a minimum processing time, i.e. $s(k)$. The SPT selection function can be seen as a specialized version of the HDF selection function, where the weights are set to either 1 or $1/t_j$. These weights correspond to metrics minimizing the average flow time or the average stretch, respectively. It can be shown that the SPT selection function is a reasonable heuristic to minimize the average response time.

[0036] D. *Earliest Deadline First (EDF)* . This selection function assigns to each job k an estimated deadline at arrival time. The deadline $d(k)$ assigned to job k is

$$d(k) = a(k) + \sqrt{s_{\max} / s_{\min}} \max_{stretch}(k) (k_{pf} \text{FRAME}_{TIME})$$

where $a(k)$ is the time of arrival of the request, s_{\max} and s_{\min} are the minimum and maximum size of a job observed so far, k_{pf} is the estimated number of frames required by the algorithm to process the job, FRAME_{TIME} is the time length of a frame. Finally, $\max_{stretch}(k)$ is the maximum stretch observed on no more than the last 10 jobs that have been completed. The EDF selection function chooses the request with the earliest deadline. It can be shown that this selection function is a reasonable online heuristic for minimizing the maximum stretch.

[0037] E. *Short Remaining Processing Time (SRPT)* . This selection function chooses the request k currently in the system having minimum k_{rpf} , where k_{rpf} is the number of frames still to be allocated to complete the request at time f . It can be shown that the SRPT selection function is a good heuristic for minimizing the average response time.

[0038] F. *Max Estimated Flow Time (MEFT)* . This selection function chooses the request k currently in the system having maximum estimation of $f - a(k) + k_{rpf}$. It can be shown that the MEFT selection function is a reasonable heuristic to minimize the maximum stretch or average response time.

[0039] G. *Max Estimated Stretch (MEST)* . This selection function chooses the request k currently in the system having maximum estimation of the stretch $(f - a(k) + k_{rpf}) / s(k)$. It can be shown that the MEST selection function is a reasonable heuristic to minimize the maximum stretch or average response time.

[0040] H. *Max Elapsed Stretch (MES)* . This selection function chooses the request k currently in the system having maximum ratio between the waiting time and demand $(f - a(k)) / s(k)$. It can be shown that the MES selection function is a reasonable heuristic to minimize the maximum stretch or average response time.

[0041]

The foregoing Detailed Description is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by the patent laws. It is

to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention. For example, the detailed description describes an embodiment of the invention with particular reference to WCDMA systems and other next-generation CDMA systems. However, the principles of the present invention could be readily extended to other wireless data network architectures, such as Wireless OFDM, etc. Such an extension could be readily implemented by one of ordinary skill in the art given the above disclosure.